


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used disk scheduling policies

Found 5,285 of 169,166

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

 Relevance scale ☐ ☐ ☐ ☐ ☐

### 1 [Special issue on the PAPA 2002 workshop: Disk scheduling policies with lookahead](#)



Alexander Thomasian, Chang Liu

 September 2002 **ACM SIGMETRICS Performance Evaluation Review**, Volume 30 Issue 2

Publisher: ACM Press

 Full text available: [pdf\(1.08 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Advances in magnetic recording technology have resulted in a rapid increase in disk capacities, but improvements in the mechanical characteristics of disks have been quite modest. For example the access time to random disk blocks has decreased by a mere factor of two, while disk capacities have increased by several orders of magnitude. High performance OLTP applications subject disks to a very demanding workload, since they require high access rates to randomly distributed disk blocks and gain I ...

**Keywords:** LOOK, SATF, SCAN, disk scheduling, scheduling policies with lookahead, simulation

### 2 [A comparative analysis of disk scheduling policies](#)



Toby J. Teorey, Tad B. Pinkerton

 March 1972 **Communications of the ACM**, Volume 15 Issue 3

Publisher: ACM Press

 Full text available: [pdf\(674.34 KB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)



Five well-known scheduling policies for movable head disks are compared using the performance criteria of expected seek time (system oriented) and expected waiting time (individual I/O request oriented). Both analytical and simulation results are obtained. The variance of waiting time is introduced as another meaningful measure of performance, showing possible discrimination against individual requests. Then the choice of a utility function to measure total performance including system orie ...

**Keywords:** access time, analytical models, auxiliary storage, direct access storage, disk analysis, disk scheduling, performance criteria, peripheral memory devices, real-time systems, response time, rotational delay, scheduling policies, seek time, simulation, storage units, time-sharing systems, waiting time



### 3 [Disk scheduling in a multimedia I/O system](#)



A. L. Narasimha Reddy, Jim Wyllie

- September 1993 **Proceedings of the first ACM international conference on Multimedia**  
**Publisher:** ACM Press  
 Full text available:  [pdf\(122.53 KB\)](#)  [ps\(242.93 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



4 Disk scheduling in a multimedia I/O system

-  A. L. N. Reddy, Jim Wyllie, K. B. R. Wijayarathne  
 February 2005 **ACM Transactions on Multimedia Computing, Communications, and Applications (TOMCCAP)**, Volume 1 Issue 1  
**Publisher:** ACM Press  
 Full text available:  [pdf\(271.67 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This article provides a retrospective of our original paper by the same title in the Proceedings of the First ACM Conference on Multimedia, published in 1993. This article examines the problem of disk scheduling in a multimedia I/O system. In a multimedia server, the disk requests may have constant data rate requirements and need guaranteed service. We propose a new scheduling algorithm, SCAN-EDF, that combines the features of SCAN type of seek optimizing algorithm with an Earliest Deadline First ...



**Keywords:** I/O systems, disk scheduling, multimedia applications, performance evaluation, real-time

5 Cello: a disk scheduling framework for next generation operating systems

-  Prashant J. Shenoy, Harrick M. Vin  
 June 1998 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1998 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems SIGMETRICS '98/PERFORMANCE '98**, Volume 26 Issue 1  
**Publisher:** ACM Press  
 Full text available:  [pdf\(1.60 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper, we present the Cello disk scheduling framework for meeting the diverse service requirements of applications. Cello employs a two-level disk scheduling architecture, consisting of a class-independent scheduler and a set of class-specific schedulers. The two levels of the framework allocate disk bandwidth at two time-scales: the class-independent scheduler governs the coarse-grain allocation of bandwidth to application classes, while the class-specific schedulers control the fine-grain ...

6 Some new disk scheduling policies and their performance

-  Alexander Thomasian, Chang Liu  
 June 2002 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '02**, Volume 30 Issue 1  
**Publisher:** ACM Press  
 Full text available:  [pdf\(123.69 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

Advances in magnetic recording technology have resulted in a rapid increase in disk capacities, but improvements in the mechanical characteristics of disks have been quite modest. For example, the access time to random disk blocks has decreased by a mere factor of two, while disk capacities have increased by several orders of magnitude. OLTP applications subject disks to a very demanding workload consisting of accesses to randomly distributed disk blocks and gain limited benefit from caching and ...

## 7 Implementation and performance of integrated application-controlled file caching, prefetching, and disk scheduling



Pei Cao, Edward W. Felten, Anna R. Karlin, Kai Li

November 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 4

**Publisher:** ACM Press

Full text available: [pdf\(609.00 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

As the performance gap between disks and micropocessors continues to increase, effective utilization of the file cache becomes increasingly important. Application-controlled file caching and prefetching can apply application-specific knowledge to improve file cache management. However, supporting application-controlled file caching and prefetching is nontrivial because caching and prefetching need to be integrated carefully, and the kernel needs to allocate cache blocks among processes ap ...

**Keywords:** application-controlled resource management, disk scheduling, file caching, file prefetching

## 8 Anticipatory scheduling: a disk scheduling framework to overcome deceptive idleness in synchronous I/O



Sitaram Iyer, Peter Druschel

October 2001 **ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM symposium on Operating systems principles SOSP '01**, Volume 35 Issue 5

**Publisher:** ACM Press

Full text available: [pdf\(1.61 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Disk schedulers in current operating systems are generally work-conserving, i.e., they schedule a request as soon as the previous request has finished. Such schedulers often require multiple outstanding requests from each process to meet system-level goals of performance and quality of service. Unfortunately, many common applications issue disk read requests in a synchronous manner, interspersing successive requests with short periods of computation. The scheduler chooses the next request too ea ...

## 9 Disk scheduling: FCFS vs.SSTF revisited



Micha Hofri

November 1980 **Communications of the ACM**, Volume 23 Issue 11

**Publisher:** ACM Press

Full text available: [pdf\(944.08 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

We report on a rather extensive simulation effort directed at evaluating the merits of two scheduling strategies, FCFS and SSTF, for moving-arm disks under stationary request arrival process. For First-Come-First-Served (FCFS) scheduling, analytic results for the mean waiting time are also given (in a closed form). If the objective of a schedule is to minimize the mean waiting time (or queue size) and its variance, the results seem to confirm the overall superiority of Shortest-Seek-Time-Fi ...

**Keywords:** FCFS, SSTF, correlated input, moving-arm disk, optimal schedule, scheduling, simulation

## 10 Performance of a two-headed disk system when serving database queries under the scan policy



Y. Manolopoulos, J. G. Kollias

September 1989 **ACM Transactions on Database Systems (TODS)**, Volume 14 Issue 3

**Publisher:** ACM PressFull text available:  [pdf\(1.12 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Disk drives with movable two-headed arms are now commercially available. The two heads are separated by a fixed number of cylinders. A major problem for optimizing disk head movement, when answering database requests, is the specification of the optimum number of cylinders separating the two heads. An earlier analytical study assumed a FCFS model and concluded that the optimum separation distance should be equal to 0.44657 of the number of cylinders  $N$  of the disk. This paper ...

**11 Performance of a mirrored disk in a real-time transaction system**

Shenze Chen, Don Towsley

April 1991 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1991 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '91**, Volume 19 Issue 1

**Publisher:** ACM PressFull text available:  [pdf\(975.06 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Disk mirroring has found widespread use in computer systems as a method for providing fault tolerance. In addition to increasing reliability, a mirrored disk can also reduce I/O response time by supporting the execution of parallel I/O requests. The improvement in I/O efficiency is extremely important in a real-time system, where each computational entity carries a deadline. In this paper, we present two classes of real-time disk scheduling policies, RT-DMQ and RT-CMQ, for a mirrored disk I/O su ...

**12 Probing the black box: Transforming policies into mechanisms with infokernel**

Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, Nathan C. Burnett, Timothy E. Denehy, Thomas J. Engle, Haryadi S. Gunawi, James A. Nugent, Florentina I. Popovici

October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**

**Publisher:** ACM PressFull text available:  [pdf\(365.12 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We describe an evolutionary path that allows operating systems to be used in a more flexible and appropriate manner by higher-level services. An infokernel exposes key pieces of information about its algorithms and internal state; thus, its default policies become mechanisms, which can be controlled from user-level. We have implemented two prototype infokernels based on the linuxtwofour and netbsdver kernels, called infolinux and infobsd, respectively. The infokernels export key abstractions as ...

**Keywords:** information, mechanism, policy**13 A simulation study of adaptive scheduling policies in interactive computer systems**

Samuel T. Chanson, Craig D. Bishop

July 1977 **ACM SIGMETRICS Performance Evaluation Review**, Volume 6 Issue 3

**Publisher:** ACM PressFull text available:  [pdf\(577.42 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#)

Recently, some work has been done in the area of dynamically adaptive scheduling in operating systems (i.e., policies that will adjust to varying workload conditions so as to maximize performance) [4],[5], [10], [11]. However, most studies deal with batch-oriented systems only. The University of British Columbia operates an IBM 370/168 running under MTS (Michigan Terminal System) which is principally used interactively. It

has been known for some time that the system is Input/Output bound. The m ...

14 A simulation study of adaptive scheduling policies in interactive computer systems

Samuel T. Chanson, Craig D. Bishop


January 1977 **Proceedings of the 9th conference on Winter simulation - Volume 2**

**Publisher:** Winter Simulation Conference

Full text available:  [pdf\(598.09 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Recently, some work has been done in the area of dynamically adaptive scheduling in operating systems (i.e., policies that will adjust to varying workload conditions so as to maximize performance) [4], [5], [10], [11]. However, most studies deal with batch-oriented systems only. The University of British Columbia operates an IBM 370/168 running under MTS (Michigan Terminal System) which is principally used interactively. It has been known for some time that the system i ...

15 Queuing and scheduling: Lexicographic QoS scheduling for parallel I/O

 Ajay Gulati, Peter Varman

July 2005 **Proceedings of the 17th annual ACM symposium on Parallelism in algorithms and architectures SPAA'05**


**Publisher:** ACM Press

Full text available:  [pdf\(261.78 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

High-end shared storage systems serving multiple independent workloads must assure that concurrently executing clients will receive a fair or agreed-upon share of system I/O resources. In a parallel I/O system an application makes requests for specific disks at different steps of its computation depending on the data layout and its computational state. Different applications contend for disk access making the problem of maintaining fair allocation challenging. We propose a model for differentiate ...

**Keywords:** QoS, fair scheduling, lexicographic minimization, parallel I/O, resource allocation, storage virtualization

16 Queueing Analysis of the Scan Policy for Moving-Head Disks


 Walter C. Oney

July 1975 **Journal of the ACM (JACM)**, Volume 22 Issue 3

**Publisher:** ACM Press

Full text available:  [pdf\(914.56 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

17 An anomaly in disk scheduling: a comparison of FCFS and SSTF seek scheduling using an empirical model for disk accesses

 Neil C. Wilhelm

January 1976 **Communications of the ACM**, Volume 19 Issue 1

**Publisher:** ACM Press

Full text available:  [pdf\(441.10 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

A model for disk accesses based on published measurements is developed. The model is used to show that under highly probable conditions, FCFS seek scheduling is superior to SSTF scheduling in the sense of having a lower mean queue length. A simple example of an arrival sequence illustrating this anomaly is presented.

**Keywords:** disk scheduling, disks, seek scheduling

## 18 An analytic behavior model for disk drives with readahead caches and request reordering



Elizabeth Shriver, Arif Merchant, John Wilkes

June 1998 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1998 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems SIGMETRICS '98/PERFORMANCE '98**, Volume 26 Issue 1

**Publisher:** ACM Press

Full text available: [pdf\(1.44 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Modern disk drives read-ahead data and reorder incoming requests in a workload-dependent fashion. This improves their performance, but makes simple analytical models of them inadequate for performance prediction, capacity planning, workload balancing, and so on. To address this problem we have developed a new analytic model for disk drives that do readahead and request reordering. We did so by developing performance models of the disk drive components (queues, caches, and the disk mechanism) and ...

## 19 Scheduling algorithms for modern disk drives



Bruce L. Worthington, Gregory R. Ganger, Yale N. Patt

May 1994 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1994 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '94**, Volume 22 Issue 1

**Publisher:** ACM Press

Full text available: [pdf\(1.01 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Disk subsystem performance can be dramatically improved by dynamically ordering, or scheduling, pending requests. Via strongly validated simulation, we examine the impact of complex logical-to-physical mappings and large prefetching caches on scheduling effectiveness. Using both synthetic workloads and traces captured from six different user environments, we arrive at three main conclusions: (1) Incorporating complex mapping information into the scheduler provides only a ma ...

## 20 A continuum of disk scheduling algorithms



Robert Geist, Stephen Daniel

January 1987 **ACM Transactions on Computer Systems (TOCS)**, Volume 5 Issue 1

**Publisher:** ACM Press

Full text available: [pdf\(866.07 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

A continuum of disk scheduling algorithms,  $V(R)$ , having endpoints  $V(0) = \text{SSTF}$  and  $V(1) = \text{SCAN}$ , is defined.  $V(R)$  maintains a current SCAN direction (in or out) and services next the request with the smallest effective distance. The effective distance of a request that lies in the current direction is its physical distance (in cylinders) from the read/write head. The effective distance of a request in the opposite direction is its physical di ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)